



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04L 12/24		A1	(11) International Publication Number: WO 97/24836
			(43) International Publication Date: 10 July 1997 (10.07.97)
<p>(21) International Application Number: PCT/FI96/00689</p> <p>(22) International Filing Date: 20 December 1996 (20.12.96)</p> <p>(30) Priority Data: 9526615.1 28 December 1995 (28.12.95) GB</p> <p>(71) Applicant (for all designated States except US): NOKIA TELECOMMUNICATIONS OY [FI/FI]; Upseerinkatu 1, FIN-02600 Espoo (FI).</p> <p>(72) Inventors; and</p> <p>(75) Inventors/Applicants (for US only): GRANT, Colin [GB/GB]; 16 Camside, Chesterton, Cambridge CB4 1PQ (GB). FAUPEL, Matthew [GB/GB]; 14 Church Lane, Milton, Cambridge CB4 6AB (GB). FRENCH, Graham [GB/GB]; 17 Cheney Way, Cambridge CB4 1UE (GB). BENDING, David [GB/GB]; 34 Chaucer Place, St Neots, Cambridgeshire PE19 3LN (GB). HEIGHES, Phil [GB/GB]; 124 The Sycamores, Milton, Cambridge CB4 6ZN (GB).</p> <p>(74) Agent: OY KOLSTER AB; Iso Roobertinkatu 23, P.O. Box 148, FIN-00121 Helsinki (FI).</p>			<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</p>

(54) Title: TELECOMMUNICATIONS NETWORK MANAGEMENT SYSTEM

(57) Abstract

A system for managing a telecommunications network, wherein the functions of the system have been divided between a single core part (CP) and at least one peripheral part (PP1-PPN) in such a way that (1) of said parts, only the core part communicates with the network elements to be managed, the core part storing all objects used by the system as master copies and maintaining information concerning said objects, (2) the peripheral part contains applications stored in the user's workstation, through which applications the services of the system are offered to the user, each peripheral part having separately stored copies of objects used by the applications of the concerned peripheral part and information indicating which applications use a particular object, and (3) the core part stores information indicating which peripheral part contains a copy of an individual object, whereby the peripheral parts and core parts communicate with each other to keep the objects in the core part and the copies in the peripheral parts consistent with each other.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

Telecommunications network management system

The present invention relates to a management system according to the preamble of attached claim 1 for managing a telecommunications network. The telecommunications network to be managed may be e.g. a SDH (Synchronous Digital Hierarchy) network, a PDH (Plesiochronous Digital Hierarchy) network, or a combination of such networks.

The basic situation in network management is usually such that an operator managing a telecommunications network, e.g. a telephone company, has a plurality of customers (i.e. network users) in addition to the physical network. The operator sells the customers various services that utilize the network. (A public network will be used herein as an example; in principle, however, the same description applies to a private operator managing e.g. an organization network). To meet customers' data transmission requirements in the physical network, the operator utilizes a number of facilities or operative processes for the provision of customer services. These operative processes can be divided into groups in accordance with the functions for which they are intended:

- Service Provisioning taking care of the performance of customer services, including e.g. invoicing customers for services.

- Operation & Maintenance for keeping the network operative to allow the usage of customer services. One of the most important functions in this respect is the supervision and repair of network faults.

- Planning & Development, the function of which is to develop network operation so as to better meet customers' needs and to increase the overall profit-

ability of the operator enterprise.

As appears from the above, network management takes place on several different levels, depending on the extent to which the functions to be performed on a specific level are associated with the overall management of the operator enterprise. The management of a telecommunications network is generally divided into four different levels, which are from bottom to top as follows:

10 - network element management layer,
 - network management layer,
 - service management layer, and
 - business management layer.

15 This division is used e.g. in the ITU-T (the former CCITT) recommendation M.3010, which specifies a kind of framework for the management architecture of a telecommunications network. The bottom layer below the above four layers is the equipment itself; these equipments are managed by installation and field 20 engineering tools.

25 The network element management layer means the management of an individual network element (such as a multiplexer or a cross-connection device) as a separate component without simultaneously paying attention to the condition of the network or other network elements. The majority of so called "network management" systems commercially available today are actually network element management systems within this 30 layer.

35 The network management layer is concerned with the management of the entire telecommunications network, such as overall management of network connections. One example is the creation of connections and the end-to-end supervision of their condition. This means that e.g. alarms detected on equipment are

not just displayed against that equipment, but they are also propagated to show what services (paths and circuits) are affected by the fault, if any. The present invention is positioned in this layer.

5 As distinct from the above, the service management layer is not concerned with technical network management. It takes care of e.g. customer data, supervision of services provided to customers, invoicing for services, and considering needs for services 10 of different types.

The business management layer is used to monitor and plan the business activities and economy of the entire enterprise, resulting in decisions affecting the lower levels.

15 At present, network management systems are changing into systems that manage the telecommunications network as a whole, whereas conventional management systems have handled only the remote control of transmission equipment, especially monitoring 20 alarms produced by the equipment. In conventional network management methods, configuration changes, such as creation of new end-to-end connections, have been laborious and time-consuming, as the end result consists of several configuration events the prerequisite 25 of which is that the maintenance staff of the network first gets an overall view of the situation and then decides on configuration changes required in individual network elements. In new network management systems, on the contrary, an overall view of the network and its condition is produced within the system, 30 and the system itself gives the required configuration commands to each transmission equipment. As a consequence, all configuration changes can be performed significantly more rapidly than previously. Such 35 developments have been accelerated by the freeing of

competition in the field of telecommunications.

The above-mentioned recommendation M.3010 specifies the management architecture as shown in Figure 1. The architecture basically consists of one or more operations systems OS connected to a data communication network DCN communicating with an actual telecommunications network which is to be managed and which includes the network elements NE managed. It is to be noted that the connections of the data communications network and those of the telecommunications network are logically distinct, although they can be implemented physically in one and the same cable. Logically, there are thus two networks: (a) a network providing services to customers, and (b) a network maintaining the service provisioning network. The management of certain transmission equipments (network elements) further requires a separate Mediation Device MD, which mainly acts as a protocol converter between a Q3 interface complying with the recommendations and transmission equipments that do not understand the protocol applied in the interface but use their own proprietary protocol. New SDH equipment, for instance, can be connected directly to the Q3 interface, whereas older PDH equipment requires a Mediation Device.

In practice, a management network for a combined SDH and PDH network may be e.g. such as shown in Figure 2. Users (network operator staff) sitting at the operation centre use network management work stations WS connected to a separate local area network WSN, which may be e.g. an Ethernet network. The management system is typically distributed in several computers of the local area network, one of the computers being a dedicated server machine SRV having a database DB containing information necessary for

managing the network. In its practical embodiment, the local area network further comprises e.g. necessary back-up devices (like DAT drives or mirrored disks) and event-logging printers (not shown).

5 The management system is connected via the above-mentioned Q3 interface e.g. to the SDH network. A variety of alternatives have been defined for the Q3 interface, so that the interface may be e.g. an X.25-type packet-switched interface or an Ethernet LAN interface. (The packet-switched interface is useful if the operator in charge of the network management also otherwise uses a packet-switched network.) In practice, control channels between the SDH network elements 21 are established in the overhead bytes of 10 the STM-N signal (N = 1, 4, 16), so that control signals between SDH equipments propagate with the payload signal (that is, also in the same physical network). Such control channels established in the overhead bytes are called Embedded Control Channels, 15 and they are formed e.g. in the STM-1 frame by the section overhead bytes D1 to D12.

20 PDH equipments, on the contrary, need manufacturer-specific solutions, wherefore they have to be connected to the management system through a separate mediation device 22.

25 The management system may also be hierarchical so that different geographical areas have their own smaller management systems that together form an integral management network. For instance, a management system covering one country may be divided geographically into smaller management systems operating in different parts of the country. Each smaller management system takes care of the management of the network portion in the concerned geographical area. In 30 the example of Figure 2, management systems MS1 and 35 MS2 are shown.

MS2 geographically apart from each other form together a single common management system and management network.

Network management standards are nowadays largely based on so-called object-oriented descriptions, though the standards do not require the use of this technique. Objects are data structures in a network management system, which describe the functions and state of a network component. An object is thus an element having certain attributes ("I am like this") and certain operations ("I can do these things"). (In the object-oriented approach, objects with the same data structure (attributes) and the same behaviour (operations) are grouped into a class. A specific implementation of an operation is called a method and each object is said to be an instance of its class.) A typical object is e.g. a cross-connection device with certain attributes (cross-connections that are active) and certain methods (e.g. make cross-connection and release cross-connection).

In a telecommunications network management system, objects can be physical ones or logical ones. Physical objects are elements that form part of the physical network. Such objects are e.g. the above-mentioned network elements (a network element is any piece of telecommunication equipment that works as a single unit to provide telecommunication functions) or physical connections (such as optical fibres, twisted pair cables, coaxial cables, radio links or satellite links). Logical objects are logical entities that do not form a single piece of the physical network. Such objects are e.g. paths and circuits. (A path is a connection of a fixed bit rate and format between two physical interfaces within the network. A circuit is a connection set up for a customer, between two physical

interfaces on the boundary of the network. Thus, a circuit usually comprises several consecutive paths.)

5 A network object may have a number of different attributes. Some attributes (such as "fault state") are used by several different types of object. In addition, for some types of network object (such as a route), it is convenient to define an attribute which consists of a collection of other attributes. Typical attributes are e.g. "availability status", "fault state" and "operational state". The attributes have different possible values, e.g. fault state can have values:

- OK. There are no problems.
- Warning. There are outstanding faults, but 15 these do not affect services.
- Degraded. Some or all of the services provided by the object are degraded.
- Failed. All the services provided by the object are lost.
- 20 - Unknown. The fault state of the object is unknown.

The "operational state", in turn, can have e.g. two different values:

- Enabled. The object can operate, either 25 completely or in part.
- Disabled. The object cannot operate at all.

As obvious from above, a network operator's biggest asset is the transmission system as such, and the management system is a tool for running efficiently a business based on the services provided by the transmission system. Therefore, the management system should be able to follow the dynamic use of the network at every moment to allow the operator to make decisions that help the operator to make the best 30 profit in all circumstances. For such decisions the 35

operator might for example want to know what should be fixed at this moment, how much spare capacity there is in the network, how the individual lines are utilized, where the restriction in the bandwidth is, etc. If the 5 network management system cannot give the operator the information needed to understand the use of the network at every moment, the operator cannot really control the network and run the business in the best possible way.

10 Most management systems are based on RPCs (RPC, Remote Procedure Call). The architecture of such a system is shown in Figure 3. The management system has typically a central machine called an RPC server. This RPC server has objects representing the various network 15 elements stored inside it. The application processes (as will also be disclosed later, an application is a functional system block containing the user interface functionality, i.e. allowing the user to configure the network and view network events 20 on the screen of the workstation) are stored in a workstation WS and they include local objects called RPC stubs which relate to the network elements to be managed. When an application wants to have information 25 on a specific object, the corresponding stub object sends a question to the remote RPC server via a link 30 between them. The RPC server then retrieves the information from the object in question, packages the result, and sends it back to the application that made the question. If the application wants to get e.g. 30 seven pieces of information to show them on the screen of the workstation, it would have to send seven messages to the RPC server and wait for seven replies. The local application itself does not know that 35 messages are transmitted to the RPC server. It is, of course, possible to send seven queries or answers at

the same time, but also in this case problems will occur if the application wants e.g. 8 or 9 pieces of information, as the structure and coding of the messages has to be changed.

5 The RPC-based systems have a central event management system EMS which broadcasts the events to all applications. If a network element NE, for instance, informs of a failure, the EMS broadcasts information about the event to all applications.

10 This kind of a system which is based on a centralized database on the RPC server works well for small systems in a normal situation. However, problems arise if one or more applications repeatedly require more information from the RPC server (for example, if there is a search going on). This can lead to a huge queue of events on the RPC server. This in turn leads, at least momentarily, to very slow applications as an individual application has to queue for every piece of information.

15 The performance of the above-described known management system is particularly problematic in emergency situations (such as a Network Element failure). In the case of emergency, all of the applications receive a fault message from the network and after that all of them request at the same time for details from the RPC server. So, when there is a major fault in the network (for example an STM-1 link is lost), the applications need a lot of information to work out the effects of the failure. It is precisely at this point where the system is suddenly overloaded and the performance of the system drops drastically. Thus, the performance of the system changes in this situation in the wrong way, as the network management system should operate in the most efficient way when there is a fault in the network.

The problem associated with the performance can be illustrated by the following example. In practice, an RPC-based system can typically handle about 100 events in one second (if there is e.g. an Ethernet LAN between the workstations and the RPC server). For example, if an STM-1 link is lost in the network to be managed, the EMS will receive about 2,000 events (an STM-1 link has about 2,000 lower-level channels). As every application is informed about the events, it may take as long as 20 seconds before information about of the fault reaches the application that needs it most. The situation is even worse if the link between the workstation and the RPC server is slower, such as a X.25 connection.

An RPC-based system may also comprise more than one RPC server. In this case, objects are spread among these servers. This, however, does not improve the performance decisively, as a specific object is nevertheless stored in a single server, and the applications have to compete with each other to gain information about the object in question.

The object of the present invention is to provide a novel network management system which does not have the above-mentioned drawbacks, that is, to provide a network management system which allows dynamic control of the network without degrading in failure situations while taking into account the operator's needs. This is achieved by means of the system according to the invention, which is characterized by what is disclosed in the characterizing portion of attached claim 1.

The idea of the present invention is to keep up a full image of the network to be managed in the system core as well as required portions of this full image outside the core in association with the

applications, and thus divide the functionality of the system in a new way. Required images are produced by copying the object data kept up in the core into required locations outside the core. What object copies are needed depends on what kind of applications are used, that is, on the copies needed by the applications.

The system according to the invention remains responsive in all circumstances in the best possible way, i.e. even when faults are detected in the network.

In the following, the invention and its preferred embodiments will be described in greater detail with reference to the example of Figures 4 to 8 of the attached drawings, in which

Figure 1 illustrates a telecommunications network management architecture;

Figure 2 shows an example of a management network for a combined SDH and PDH network,

Figure 3 schematically illustrates the functional architecture of a known network management system;

Figure 4 illustrates the functional architecture of the management system according to the present invention in its basic form;

Figure 5 illustrates the function of the system according to Figure 4 when a fault is received from the network to be managed;

Figure 6 illustrates a preferred functional architecture of the management system according to the present invention;

Figures 7 and 8 illustrate the process used in the system to keep the local copies consistent with the master copies, and vice versa.

The functional architecture of the management

system according to the present invention is illustrated in Figure 4. A system with this kind of architecture is able to operate e.g. in the system shown in Figure 2. The system comprises a single core part CP and one or more peripheral parts PP1...PPN, which are connected to the core part of the system. In this example, each peripheral part is intended for one user (operator/employee) of the system.

The system typically has a dedicated core computer, such as a server SRV (also shown in Figure 2). The core part of the system keeps up an updated model of the state and operation of the telecommunications network TN to be managed, of which a single network element is shown in the figure. As described earlier, the core part uses objects to keep up this model. The core computer SRV is connected to one or several workstations WS1...WSN each including one peripheral part. The functional processes of the system, which run on the core computer and on the workstations, have been marked with ovals. These processes will be described in the following.

An application (indicated with the references APPLN1...APPLN4 in Figure 4) is a functional system block which allows the user to configure the network and view network events on the screen of the workstation. The application is thus a system part offering the system's services to the user. An individual workstation is therefore also called an application server.

A Session Handler SH is a functional system part that handles a single user's session with the system and forwards messages between the core and the user's applications. The Session Handler can be associated with the initial system log-on window through which a user logs on and establishes a session with the core.

5 The Session Handler starts up various applications according to commands given by the user. When the user logs out, the Session Handler performs the necessary closing-down and deregistering operations to terminate the session (the term registering will be explained later).

10 A Core Access Process CAP verifies that the Session Handler belongs to a known user on a known host. Having done so, it acts as the gateway to the system core. For each CAP in the core part there is a corresponding Session Handler in the peripheral part.

15 A Database Process DP controls access to the database DB to store persistent information (master objects of the management system). The main functionality of this process is:

- to convert persistent objects to and from a storable format,
- to store the current state of all persistent objects in the database DB and to keep that database up to date by the use of change and event messages sent to it by a Modelling Process MP of the core part,
- to handle the forwarding of event information to interested parts of the system.

25 The Modelling Process MP is responsible for maintaining the currency of the core's model of the network to be managed and dealing with changes to it. The main functions of the Modelling Process are:

- to accept change and event indications from either EIP or CAP (i.e. either from the network or from the user) and validate them,
- to apply changes to the model if they are valid and to determine the results of these changes,
- to pass the changes to an appropriate External Interface Process EIP, if the changes require corresponding changes in the network to be managed,

- to generate events based on the changes and events received (e.g. fault events on paths and circuits).

5. The External Interface Process EIP is a functional system part that converts data from the external world (data from the network to be managed) into the internal world of the management system. As far as the system user is concerned, he or she just sees a network element NE, like a multiplexer or a 10 cross-connect device, without having to know the manufacturer or the version of the device. Different types of External Interface Processes are used for different classes of equipment. The main functionality of each EIP is:

15 - to monitor the network for events and to exchange notifications with the network,

- to translate these events and notifications into equivalent events that are applied to the Modelling Process MP,

20 - to pass these events and notifications to the Modelling Process,

25 - to accept changes received from the Modelling Process; translate them into the format of the network model; and send them to the network and then inform the Modelling Process when all commands relating to each original change have succeeded, failed or timed out.

30 In fact, the Interface Processes could also be included in the core part, but Figure 4 shows them separate, as they form the interface of the management system with the external world. In practice, communication with the external world can take place e.g. through a 7-layer OSI (Open Systems Interface) protocol stack defined for the Q3 protocol. The 35 interface processes are described in the Applicant's

parallel application GB-9526614.4, which is referred to for a more detailed description.

5 A Fault Propagation Agent Process FPAP is a functional system part that works out the actual impact of a received event in the system. In order to be able to do this the FPAP uses information about the object class in question. This information it receives from the Database Process DP.

10 As the network elements cannot provide all the information needed by the applications, the database DB must store information e.g. about individual network elements and interconnections and interrelations between the network elements, the operations the network elements are capable of performing, and the 15 services provided by the network elements. This information is stored in the form of an object model representing the transmission network TN in terms of objects and their attributes and methods performed on these attributes. Thus, an object is the representation, in the object model, of one of the resources to 20 be managed.

25 Applications use these objects that form the image of the network to be managed within the system. Therefore, they have to get objects from the core part of the system. An application can also create a new core object to database DB, e.g. as a result of a user command.

30 Applications also register to receive various events. The registration is passed to the Session Handler and further to the core part of the system. The registration is stored in the form of a filter F1 in association with the Database Process DP. This filter finds out on the basis of the event type, to which Core Access Processes and to which applications 35 a certain type of event is to be transmitted (i.e.

what users are interested in a certain type of event). Events from the core of the system are duplicated in the Session Handler and sent to all applications which have registered to receive them. In this way, all of 5 the user's applications can receive the event at the same time, so that information presented to the user by the different applications is consistent.

In this way the core part of the system knows to which application a certain event has to be transmitted. This also allows unwanted events to be 10 filtered out as soon as possible (in the core) and reduces the bandwidth used between applications and the core part of the system. The core part is thus able to filter events up to the application level. It 15 is also possible to use filters in steps so that the core part knows only up to the level of the Session Handler what event has to be transmitted to the peripheral part. The Session Handler thereafter has a new filter, which indicates in which event a specific 20 application inside the concerned Session Handler is interested. The first alternative, however, is to be preferred in that it does not make the peripheral part (Session Handler) too complicated.

The use of filters is described in the 25 Applicants's parallel patent application GB-9526600.3, which is referred to for a more detailed description.

According to the present invention, part of the objects stored in the core part of the management system are placed in a cache CH at each Session 30 Handler SH. (Although a cache is a certain kind of memory, in this connection caching means only that a copy of something is kept to get fast access to it.) Thus, the Session Handler has a copy of any object an application has fetched from the core part of the 35 system. The Session Handler thus also stores object-

specific information indicating which one of its applications uses a specific object. This information can be stored e.g. in the form of a table TBL1 shown in Figure 4. The core part correspondingly knows which 5 Session Handler has a copy of a specific master object. This information can be stored in the form of a table TBL2 in association with the Database Process,

At the system start-up, the Session Handler has no copies of objects. When the application needs a 10 copy for the first time, the copy is fetched from the core part. As the transmitted retrieve message includes information specifying the object to be retrieved and the application that needs the copy, the Session Handler is able to keep up a table (TBL1) 15 indicating which application uses a specific object, and the Database Process is able to keep up a table indicating which Session Handler uses a copy of a specific master object.

If an application fetches an object of which the 20 Session Handler already has a copy, the application receives a copy of the object from the Session Handler without having to interact with the core part of the system. When the application ceases using an object, it transmits a message as an indication of this to the 25 core part through the Session Handler, as a result of which the copy is removed from the memory of the Session Handler and the tables are updated accordingly. If the object is needed again, the copy is fetched from the core in the same way as at the 30 system start-up.

When changes are made to an object in an application, these changes are transmitted through the Session Handler to the core, which reapplies the 35 changes to the master objects to keep the state of the master copies consistent with the state of the copies

in the peripheral part. Once the changes have been accepted and validated by the MP and stored in the DB they are passed onto all other interested Session Handlers and applications so that their copies of the objects are updated.

In the following the operation of the system according to the invention will be described more fully with reference to Figure 5 and using an alarm from the network as an example. For the sake of clarity, filter F1 and table TBL1 are not shown in Figure 5.

At first, a fault occurs in a network element NE of the network to be managed (a single network element is shown in the figure). The network element thus sends an alarm message F concerning the fault to the respective EIP. The EIP first of all retrieves the object that represents the particular network element from the database (arrow 1) and applies the change to that network element. Depending on the location of the fault, it may also be necessary for the EIP to retrieve objects associated with the first object (arrow 2). The fault indication may come e.g. from a 4 x 2 Mbit/s multiplexer, but the EIP finds out on the basis of the fault data and the multiplexer object it has retrieved that the actual fault is e.g. in one of its 2-Mbit/s input interfaces. The EIP thus finds out on the basis of the information it receives from the network which one of the system objects is concerned (e.g. the above-mentioned interface object) and what the fault is like.

The EIP then sends a notification of the object and the change occurred in it to the Modelling Process. This message is indicated with the reference Δ, which indicates that only information about the change that has occurred in the object is sent to the

Modelling Process.

5 The Modelling Process thereafter gets a copy of that object (arrow 3) from the database, applies the alarm to it, and sends it back as a delta (Δ) change message (arrow 4), which describes the changes of that (interface) object. At this stage the change has thus been updated in a master copy of the concerned object in the database DB. The fault code F received from the network to be managed can also indicate more than one 10 alarm, which are then similarly updated in the database. According to a preferred embodiment of the invention, the Modelling Process can also implement caching of objects so that frequently changed objects do not need to be repeatedly fetched from the database 15 DB.

15 The Core Access Process CAP and the Fault Propagation Agent Process FPAP have previously registered to the Database Process DP to inform the Database Process what type of events they are interested in receiving. As one of those types is a 20 new fault, the delta change is then copied to the FPAP process (arrow 5).

25 The purpose of the FPAP process is, as mentioned above, to work out the actual impact of the fault within the management system. If e.g. the 2-Mbit/s interface is connected to a 2-Mbit/s path object which has e.g. up to 31 times 64 kbit/s circuit objects, the FPAP process will send a report (arrow 6) to the Modelling Process telling the Modelling Process MP to 30 change the state of all these objects to be "failed". On the basis of the received information the Modelling Process retrieves all these objects from the database (arrow 7), changes them, and stores them again in the database DB (arrow 8). All of the related master 35 objects in the database DB have now been updated.

At the same time as the delta change is copied into the FPAP, it is also copied into access processes CAP (arrow 5') that have registered to the Database Process as processes that want to receive an event of that type. Accordingly, the delta change is sent once from the Core Access Process in question to the corresponding Session Handler (arrow 6'). It may then be distributed n times between the applications (arrows 7').

Consequently, the Database Process has sent one message to the FPAP and one message to each CAP that wants to receive any event of that type. In the Session Handler the received message has been copied n times into all interested applications. Depending on the application in question, some of them (like application 1 in Figure 5) may then respond to the received message by asking for more information (arrow 8'). This message is then transmitted to the Database Process (arrows 9' and 10') and a reply is sent back (arrows 11', 12' and 13').

For some of the applications, the delta change (arrow 7') may not be sufficient but the applications may also need the object to get more information on the actual impact of the delta change, for example. The applications are event-based. This means that when they receive an event (delta change), they act according to predetermined rules, like "get the object, process the object this way, present it on the screen of the workstation, apply it to the printer, and put it into a log".

It is further to be mentioned that the Object Model of the system stores information concerning the way in which information on a fault of each particular object class is forwarded. As a consequence, the Database Process DP knows that e.g. in the case of the

2-Mbit/s interface object, fault information has to be forwarded to all objects which this interface object provides service for. However, if the fault is e.g. on one of the interfaces of a change-over unit providing protection at some level, the error is applied only to that interface (and not any further), as the system knows that this object has a back-up route, i.e. the concerned path and circuits are not out of service. In this way the basic behaviour of the system in fault conditions can be modelled inside the management system.

In the above example, each user has a Session Handler, and a separate (logical) link to the core part is arranged for each Session Handler/CAP pair. If there is more than one user on the same workstation, the same delta change or the same object has to be copied several times (once for each user), i.e. the caching of objects will be done several times. This is not the most efficient way. In particular, if the link between the Core Access Process and the Session Handler is slow, the need to send the same information twice or even more times deteriorates the performance. This drawback can be dispensed with by the embodiment shown in Figure 6, where the workstation has one common Workstation Session Handler WSH, which is connected to the different user Session Handlers SH. In this case, there is thus a single logical link between one workstation WS and the core part of the system, irrespective of the number of users of the workstation.

The peripheral part PP in the workstation comprises several user parts; the example shown in Figure 6 has two user parts UP1 and UP2 (each user part comprises the applications of the user, the Session Handler SH of the user, and the corresponding

cache CH. The caching of objects in the workstation takes place cumulatively on several different levels (here on two levels) so that, in each user Session Handler cache CH, copies of objects needed by the applications connected to that Session Handler are cached, and in the Workstation Session Handler copies of all objects needed by the applications of the workstation. The Workstation Session Handler thus knows which objects each Session Handler stores. If one of the applications needs an object which it does not currently store in its memory space, it first asks its own Session Handler whether it has the object. If this does not have a copy, it asks the Workstation Session Handler. Accordingly, the application may get the object either from another application inside the same user session or from another user's application inside the workstation session. Only in cases where not even the Workstation Session Handler has a copy of the object, the core part has to be asked for it. In this way, traffic between the workstation and the core server will be minimized, which is important particularly when the workstation and core server have been interconnected by a slow link (for instance, when they are remote from each other e.g. through a slow packet-switched network). In this embodiment, correspondingly, the table TBL2 of the core part contains information indicating which one of the Workstation Session Handlers has a copy of each particular object.

In the same way as in the first embodiment, the applications can register to the core so as to receive various events so that the core part can filter (using filter F1) events up to the application level. Alternatively, filtering can take place in steps so that filter F1 contains information indicating which

type of events each particular workstation is interested in. Then, an additional filter F2 (or a search table) must be provided in association with the Workstation Session Handler of each workstation. The 5 filter F2 indicates on the basis of the event type, which Session Handler (i.e. which user) wants to get that particular event.

In order that the local copies and master objects could be kept consistent in the above-described system, it is to be noted that several 10 events may simultaneously try to make changes that are inconsistent with each other. In the following, a procedure used in the system according to the invention to solve this problem will be described with 15 reference to Figures 7 and 8 by using an application that creates circuits between objects as an example. If the user of the system wants to set up a circuit between two network elements, say A and B, he or she gives a corresponding command from his workstation. 20 Then the application in charge of setting up circuits (Circuit Manager in figure 7) creates a link between objects A and B representing the network elements A and B. This action is recorded in the application, as shown in Figure 8. (This is what is known as action 25 recording. It is a record of things that the user is trying to do, looking it at an application level.) This action then results in a set of changes (like set attribute X of A to be Y, set attribute X of B to be Y, set direction..., etc.) according to predetermined 30 rules. These individual changes on individual parts of each object are then recorded, which is shown as the Changes Recording stage in Figure 8.

In a normal case, the application transmits 35 these actions through the Session Handler to the core part, and the core part reapplies the actions to the

5. master objects, in order that the information of the core part on the objects would be consistent with the information of the application. The core part further transmits information on the resulting changes to all processes that have notified that they are interested in receiving such events.

10 When the application makes changes to the object, the core part may, however, receive information from the network to be managed indicating that there is a failure in the network element B, for instance (which may mean e.g. that the link between objects A and B cannot be created). The core part thus immediately transmit a message to the application as an indication of the failure; in the figure, the message is indicated with the reference AlarmEvAVC (Attribute Value Change event). This message carries a description of a change to the attribute value in question (it is thus a message of the same kind as the delta message in Figure 5). In this case the received information means that we now have a new alarm in network element B. The application then proceeds through the following steps:

25 1. All the changes that have been made at the Changes Recording stage are reversed, i.e. the application is returned to the same state where it was before the circuit set up command was given by the user.

30 2. Changes from outside the application are applied to the objects. In this example, this means that object B has to be moved into fault state by using the information contained in the AlarmEvAVC message.

35 At this stage the state of the objects (in this example object B) matches those in the core part of the system.

3. The Actions recorded earlier are reapplied one by one, and if any action fails, the application is informed. This means that the changes that the application made are retested to see whether they can 5 still be made. This is because it is possible that the change that has occurred in the network was so insignificant that it is still possible to set up a circuit between network elements A and B.

10 The above-described procedure ensures that the changes made are always done on top of the core (as the changes made in the application are first taken away and then the changes that have been done in the core are added).

15 Even though the invention has been described above referring to the example shown in the attached drawing, it is obvious that the invention is not restricted to the example, but it can be modified in many ways within the scope of the inventive idea presented above and in the attached claims. As mentioned above, the management system includes a variety 20 of objects, representing both logical and physical entities. In the attached claims, all these objects are set forth as "the objects that relate to the network elements to be managed", i.e. objects that represent the environment to be managed. Even though, 25 in the typical example set forth above, the core part and peripheral parts are on different computers, so that they may be geographical very far from each other, it is equally possible to use the entire system on one and the same computer (which is typically a UNIX computer).

30 The above mentioned parallel applications 9526614.4 and 9526600.3 are hereby incorporated by reference.

Claims:

1. A system for managing a telecommunications network, said telecommunications network comprising several network elements to be managed by the system,

5 said system comprising a management centre having at least one workstation accomplishing a man-machine interface and allowing a manager to control the system, and the system being capable of providing the manager with information on the network, said at least one workstation being connected to a database including information about the network to be managed, said information being in the form of objects that relate to the network elements to be managed and in the form of references between the different objects,

10 said references indicating the dependencies between the objects, whereby said management centre is connected to said network elements by data communication links such that the manager can initiate an operation on and receive information from a network item to be managed, characterized in that functions of the system have been divided between a single core part and at least one peripheral part in such a way that

15 - of said parts, only the core part communicates with the network elements to be managed, the core part storing all objects used by the system as master copies and maintaining information concerning said objects,

20 - the peripheral part contains applications stored in the user's workstation, through which applications the services of the system are offered to the user, each peripheral part having separately stored copies of objects used by the applications of the concerned peripheral part and information indicating which applications use a particular object,

25

30

35

and

5 - the core part stores information indicating which peripheral part contains a copy of an individual object, whereby the peripheral parts and core parts communicate with each other to keep the objects in the core part and the copies in the peripheral parts consistent with each other.

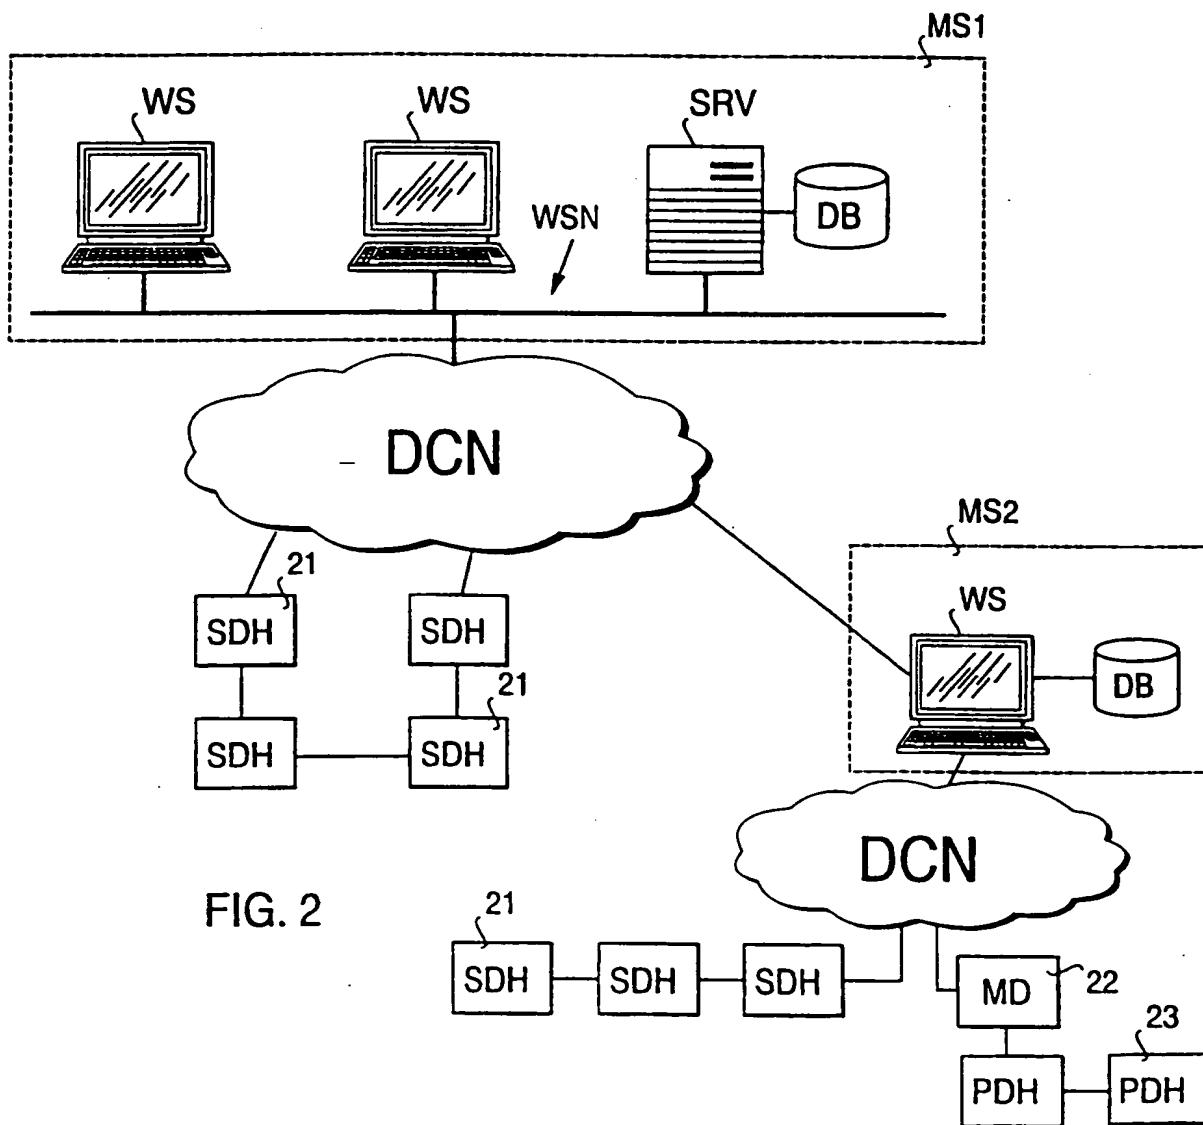
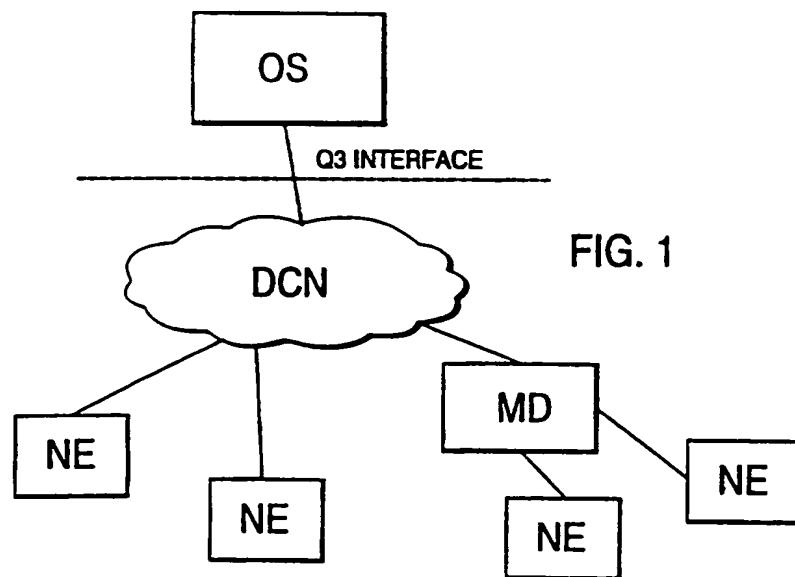
10 2. A system according to claim 1, characterized in that a copy of an object is removed from the peripheral part when none of the applications needs it any longer, whereby the information indicating which applications use a particular object is updated each time an application retrieves a new object copy from the core part.

15 3. A system according to claim 1, characterized in that one peripheral part contains applications needed by more than one user, the peripheral part storing copies of objects cumulatively so that the objects needed by each user's applications are stored separately in their own user-specific groups, in addition to which objects needed by all applications of the peripheral part are stored in their own group.

25 4. A system according to claim 1, characterized in that each peripheral part is in its own workstation and each core part in its own server machine, and that there is a single link between the peripheral part and the core part for forwarding messages between the peripheral part and core part.

30 35 5. A system according to claim 1, characterized in that the core part further stores information indicating to which peripheral part different type of events from the network are to be sent.

6. A system for managing a telecommunications network, substantially as hereinbefore described with reference to Figs 4 to 8 of the accompanying drawings.



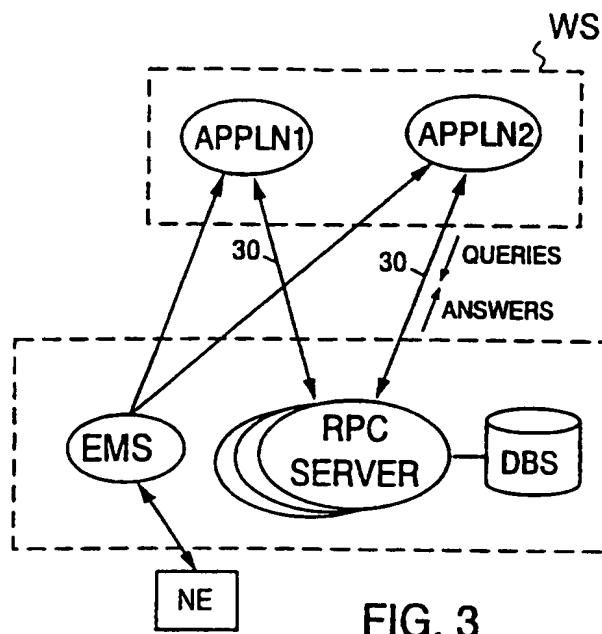


FIG. 3

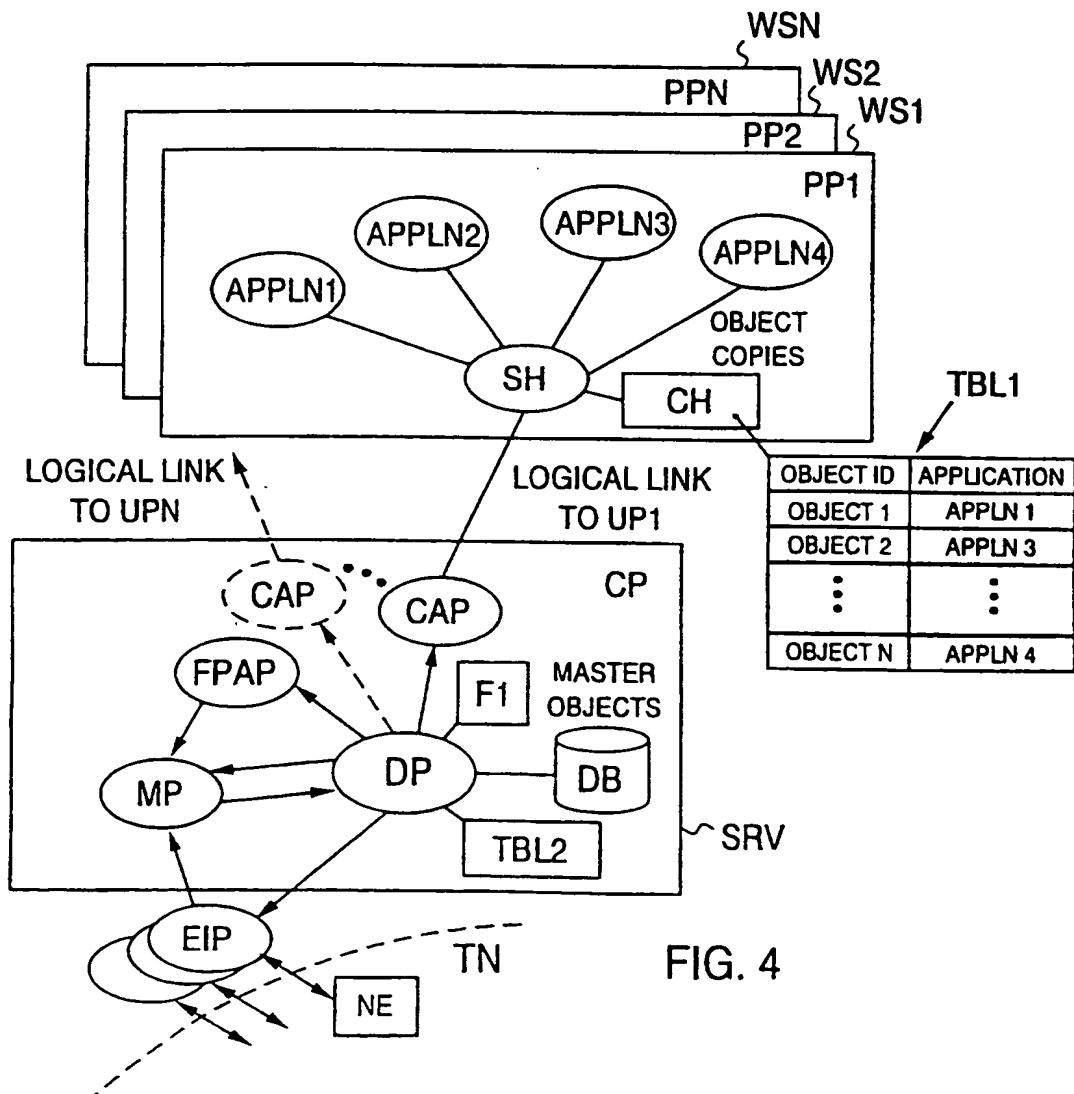
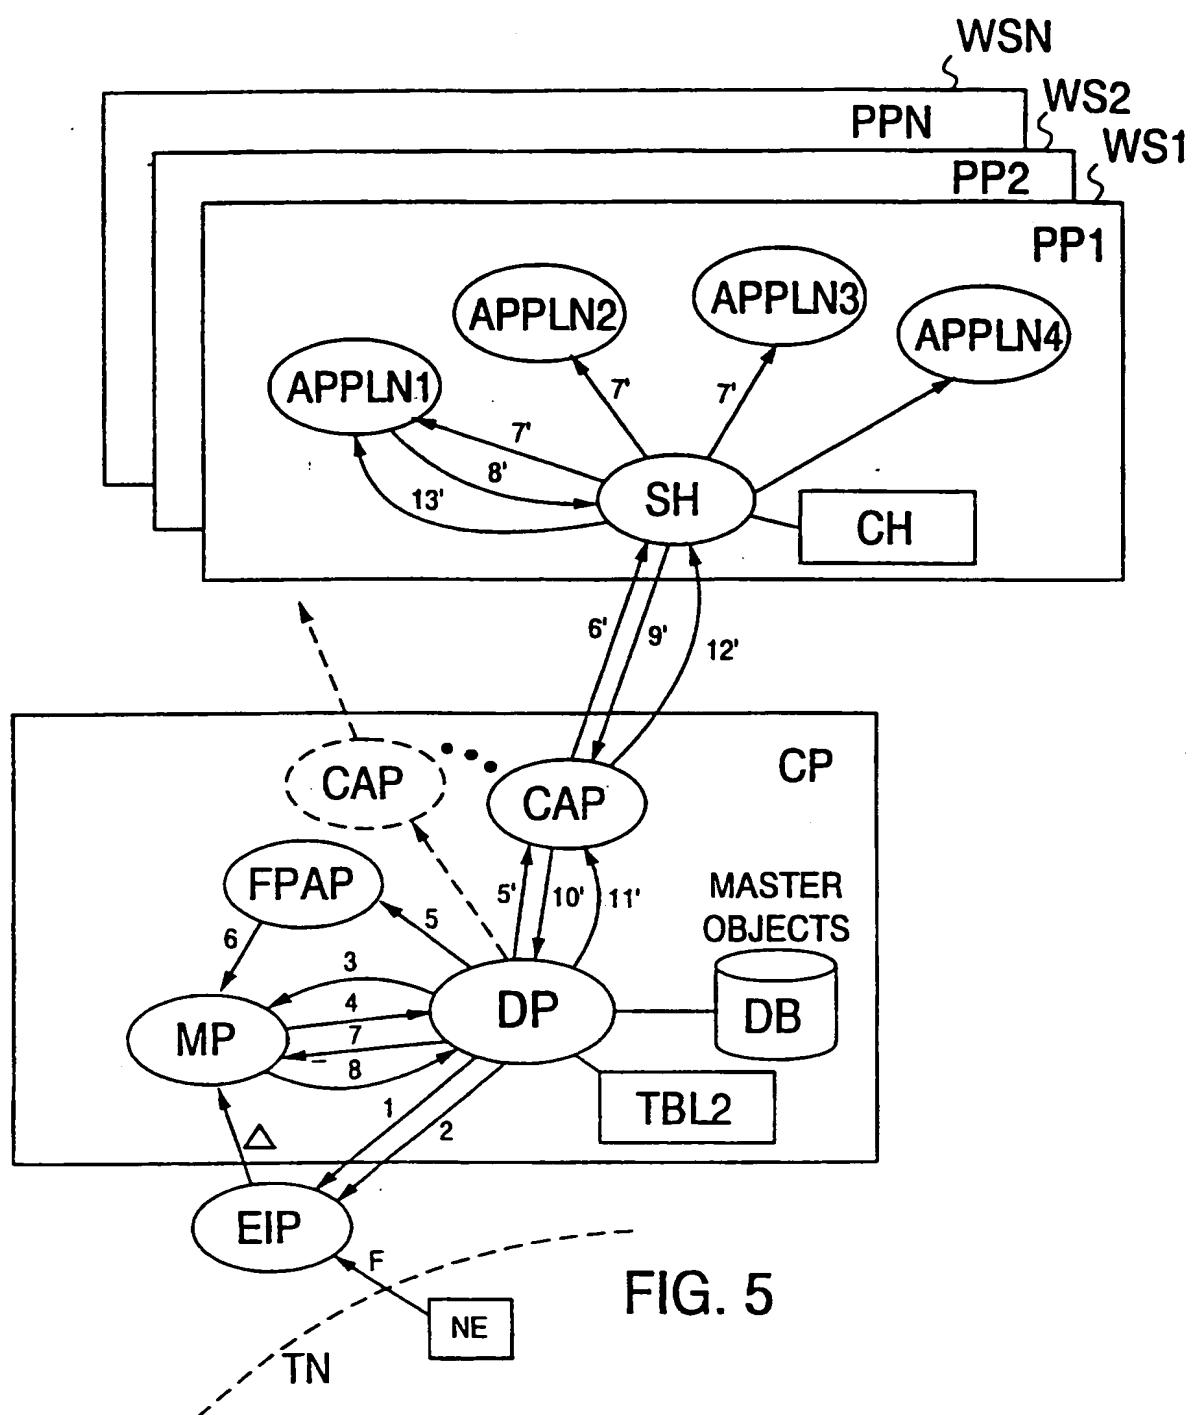
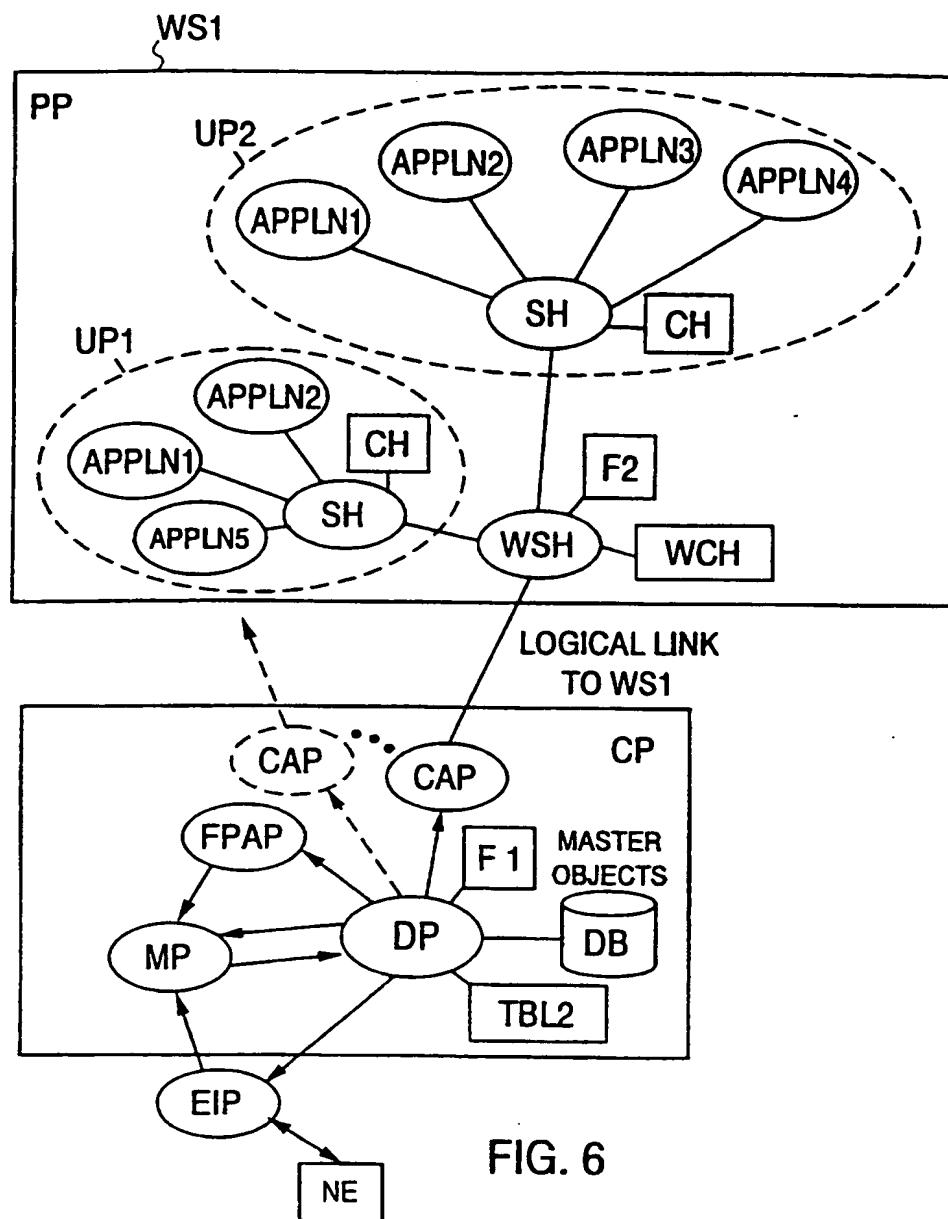


FIG. 4





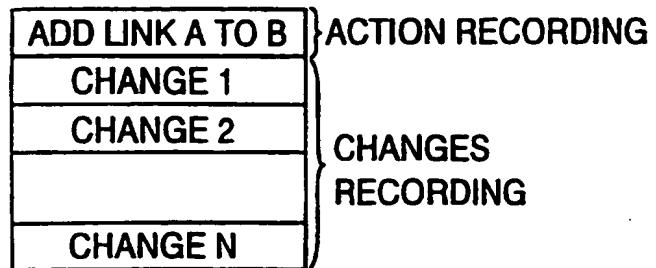


FIG. 8

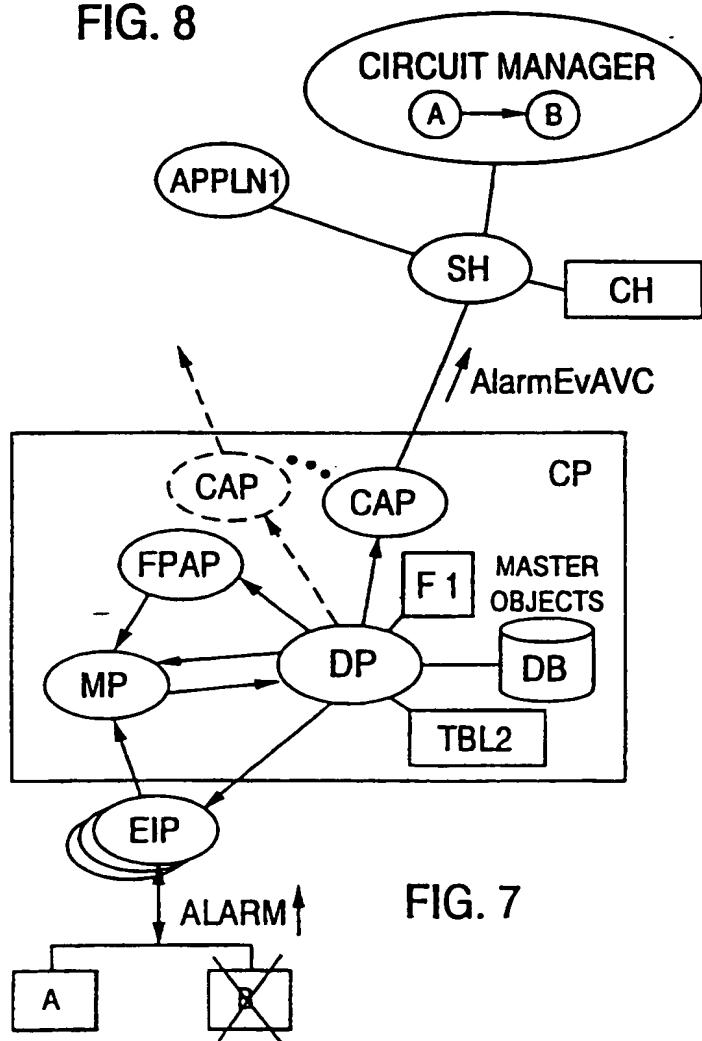


FIG. 7

1
INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 96/00689

A. CLASSIFICATION OF SUBJECT MATTER

IPC6: H04L 12/24

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPI, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	JP 7273764 A (FUJITSU LTD), 20 October 1995 (20.10.95) --	1-5
P,X	WO 9638950 A1 (DSC COMMUNICATIONS CORPORATION), 5 December 1996 (05.12.96), page 18, line 26 - line 32; page 22, line 22 - page 24, line 14 --	1-5
A	Data Communications, Volume 24, No 1, January 1995, (USA), M.Jander, "Network Management and Security" page 116 - page 118 -- -----	1-5

 Further documents are listed in the continuation of Box C. See patent family annex.

- * Special categories of cited documents:
- "A" document defining the general state of the art which is not considered to be of particular relevance
- "B" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search

25 April 1997

Date of mailing of the international search report

29 -04- 1997

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. + 46 8 666 02 86

Authorized officer

Anders Ströbeck
Telephone No. + 46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI96/00689

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.: because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.: 6 because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
Claim 6 is not clear and concise (see PCT Article 6). It relies on figures 4 - 8 and, hence, does not comply with rule 6.2(a) to such an extent that this International Searching Authority considers that no meaningful search can be performed (see Article 17.2(a) (ii)).
3. Claims Nos.: because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

The additional search fees were accompanied by the applicant's protest.

No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/FI 96/00689

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
JP 7273764 A	20/10/95	NONE	
-----	-----	-----	-----
WO 9638950 A1	05/12/96	AU 5984996 A	18/12/96
		AU 6034896 A	18/12/96
		AU 6037596 A	18/12/96
		AU 6037796 A	18/12/96
		AU 6044396 A	18/12/96
		AU 6148896 A	18/12/96
		GB 2301735 A	11/12/96
		GB 2301745 A	11/12/96
		GB 2301746 A	11/12/96
		GB 2301754 A	11/12/96
		GB 9511192 D	00/00/00
		GB 9603781 D	00/00/00
		GB 9603782 D	00/00/00
		GB 9603787 D	00/00/00
		WO 9638930 A	05/12/96
		WO 9638953 A	05/12/96
		WO 9638954 A	05/12/96
		WO 9638965 A	05/12/96
		WO 9638967 A	05/12/96
		GB 9619137 D	00/00/00
		GB 9619163 D	00/00/00
		GB 9619226 D	00/00/00
-----	-----	-----	-----

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

This Page Blank (uspto)